



06-02-05

IFW/

PATENT
Attorney Docket No.: 16869B-074700US
Client Ref. No.: HAL260
(340300744US01)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

KENJI YAMAGAMI

Application No.: 10/602,223

Filed: June 23, 2003

For: THREE DATA CENTER
REMOTE COPY SYSTEM

Customer No.: 20350

Examiner: Unassigned

Technology Center/Art Unit: 2182

Confirmation No.: 1094

**PETITION TO MAKE SPECIAL FOR
NEW APPLICATION UNDER M.P.E.P.
§ 708.02, VIII & 37 C.F.R. § 1.102(d)**

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is a petition to make special the above-identified application under MPEP § 708.02, VIII & 37 C.F.R. § 1.102(d). The application has not received any examination by an Examiner.

(a) The Commissioner is authorized to charge the petition fee of \$130 under 37 C.F.R. § 1.17(i) and any other fees associated with this paper to Deposit Account 20-1430.

(b) All the claims are believed to be directed to a single invention. If the Office determines that all the claims presented are not obviously directed to a single invention, then Applicants will make an election without traverse as a prerequisite to the grant of special status.

(c) Pre-examination searches were made of U.S. issued patents, including a classification search and a foreign patent database search. The searches were performed on or around March 11, 2005, and were conducted by a professional search firm, Mattingly, Stanger Malur & Brundidge, P.C. The classification search covered Class 707 (subclasses 202 and 204), Class 709 (subclasses 217, 218, and 219), and Class 711 (subclasses 161 and 162). Because of the large size of these subclasses, keywords were used to narrow of number of documents returned. The foreign patent database search was conducted using Espacenet database and Japanese patent database.

(d) The following references, copies of which are attached herewith, are deemed most closely related to the subject matter encompassed by the claims:

- (1) U.S. Patent No. 5,937,414;
- (2) U.S. Patent No. 6,304,980 B1;
- (3) U.S. Patent No. 6,647,474 B2;
- (4) U.S. Patent No. 6,684,306 B1;
- (5) U.S. Patent No. 6,823,349 B1;
- (6) U.S. Patent Publication No. 2003/0051111 A1;
- (7) U.S. Patent Publication No. 2004/0260902 A1;
- (8) U.S. Patent No. 6,209,003 B1;
- (9) U.S. Patent Publication No. 2003/0177321 A1; and
- (10) U.S. Patent Application No. 10/832,839, filed April 26, 2004.

(e) Set forth below is a detailed discussion of references which points out with particularity how the claimed subject matter is distinguishable over the references.

A. Claimed Embodiments of the Present Invention

The claimed embodiments relate to a remote copy system configured to perform a remote copy function involving at least three storage systems.

Independent claim 1 recites a remote copy system, comprising a first storage system including a first storage controller and a first data volume, the first storage controller being configured to control data access requests to the first data volume, the first storage system being configured to store write data in the first data volume upon receiving a write request from a first host associated with the first storage system and generate a journal including control data and journal data; a second storage system including a second data volume, the second storage system being configured to receive and store data corresponding to the write data to the first data volume in the second data volume; and a third storage system including a third data volume, the third storage system being configured to receive and store data corresponding to the write data to the first data volume in the third data volume. The write request from the first host to the first storage system completes after the write data are copied to the second storage system. The journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data.

Independent claim 26 recites a remote copy system, comprising a first storage system including a first storage controller and a first data volume, the first storage controller being configured to control data access requests to the first data volume, the first storage system being configured to store write data in the first data volume upon receiving a write request from a first host associated with the first storage system and generate a journal including control data and journal data; a second storage system including a second data volume, the second storage system being configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and a third storage system including a third data volume, the third storage system being configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data.

One of the benefits that may be derived is a remote copy system that provides the benefits of the synchronous and asynchronous modes based on journal data of a journal generated in the first storage system.

B. Discussion of the References

1. U.S. Patent No. 5,937,414

The patent to Souder et al., US 5937414, discloses the replication of data to some destination sites in a synchronous manner, and replicating that same data to other destination sites in an asynchronous manner. This allow synchronous updating to selected remote copies of replicated data where data integrity is a high priority, and asynchronous propagation to remaining copies of replicated data in order to allow transactions to be committed locally regardless of whether the transaction is committed at a remote copy of the replicated data.

The reference is directed to a mixed propagation configuration in which the propagation mode (synchronous or asynchronous) is selectable by a user on a site-to-site basis for each replicated body of data to provide a mix of availability and consistency. It does not, however, teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

2. U.S. Patent No. 6,304,980 B1

The patent to Beardsley et al., US 6304980, discloses a data backup system with a primary and redundant backup storage, where the system automatically switches to the mirroring backup storage when an error occurs at the primary storage. The system features a

synchronous remote copy feature that sends primary data to the secondary location and confirms the reception of the data before ending the primary I/O operation. The system also features an asynchronous remote copy feature in which the primary data is sent to the secondary site which might be thousands of miles away. See, e.g., Abstract; column 3, lines 33-46; and column 4, lines 8-22.

The reference is directed to peer-to-peer backup system between a primary site and a secondary site. It does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

3. U.S. Patent No. 6,647,474 B2

The patent to Yanai et al., US 6647474, discloses a data storage system that provides for the copying of data from a primary data storage system to a physically-remote secondary data storage system. The system may operate in either a synchronous mode, or an asynchronous mode. See, e.g., Abstract; and column 6, lines 17-39. The system also has primary and secondary data storage system controllers that coordinate the copying of primary data to the secondary data storage system and maintaining a list of primary data which is to be copied to the secondary storage device. The secondary data storage system controller provides an acknowledgement to the primary data storage system controller that the primary data to be copied to the secondary data storage system has been received. See, e.g., Abstract; and column 2, lines 26-57.

The reference relates to remote data mirroring between a primary data storage system and a remote secondary storage system. It does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first

storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

4. U.S. Patent No. 6,684,306 B1

The patent to Nagasawa et al., US 6684306, discloses a data storage system that can switch between an asynchronous copying mode to a synchronous copying mode to copy the data from a first storage system to a second storage system. This switching occurs on an impending hazardous event. See, e.g., Abstract; and column 2, lines 3-10.

The reference is directed to switching copying mode for data backup after the recognition of a probable occurrence of a hazardous event. It does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

5. U.S. Patent No. 6,823,349 B1

The patent to Taylor et al., US 6823349, discloses a method and system in which a host computer writes data to a primary storage site, and then synchronously writes the data to a secondary storage site. If a failure or fault occurs when writing to a secondary

storage site, a fracture log is created that represents the changed regions in the primary data at the primary storage site in response to the write request. This log can be used to synchronize the secondary storage site to the primary storage site once writing to the secondary site is restored (see Abstract).

The reference discloses the use of a fracture log, which is a bitmap representative of changed regions in the primary image at the primary image site in response to the write request, to synchronize the primary image and the secondary image when writing to the secondary image is restored. The fracture log is not a journal that includes the control data and journal data as recited in the claims. Thus, the reference fails to teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. It also fails to teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

6. U.S. Patent Publication No. 2003/0051111 A1

The published patent application to Nakano et al., US 20030051111, discloses three data centers that comprise a storage sub-system. The first data center can synchronously copy data to a second data center, and the first data center can also asynchronously copy data to a third data center. Each storage sub-system includes a function where, during normal operation, data can be exchanged and the data update state can be obtained by the storage sub-systems located in the two data centers that do not directly engage in data transmission. See, e.g., Abstract and paragraph [0079].

The reference discloses three data centers with synchronous copying from the first data center to the second data center and asynchronous copying from the first data center to the third data center, but it fails to teach the use of a journal in the first storage system to

store control data and journal data and the use of the journal data in the third data volume. Data is stored and managed in cache. See paragraphs [0098]-[0106], [0146], [0194], and [0264]-[0269]. The reference does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

7. U.S. Patent Publication No. 2004/0260902 A1

The patent to Stanley et al., US 20040260902, discloses a method and system for the remote copying of data, and it involves data being copied synchronously from a first storage unit to a second storage unit, and then the data is copied asynchronously from the second storage unit to the third storage unit. See, e.g., Abstract and paragraphs [0011]-[0013]. The reference also discloses that the first storage unit, second storage unit, and third storage unit each contain a first storage volume in a first storage control unit. The second control unit is in a location that is within synchronous communication with the first control unit. The third storage control unit is in a location that is beyond synchronous communication from the first storage control unit, and therefore data is copied there asynchronously from the second storage unit. See, e.g., paragraphs [0011]-[0012].

The reference discloses three storage units with synchronous copying from the first storage unit to the second storage unit, but asynchronous copying from the second storage unit to the third storage unit. Thus, it does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured

to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

8. U.S. Patent No. 6,209,003 B1

The patent to Gagne et al., US 6209902, discloses a local site 21, a first remote site 22, and a second remote site 23, which allow cascading data through redundant data storage units. Data is transferred from the local site 21 to the first remote site 22, and from the first remote site 22 to the second remote site 23.

The reference is directed to a three data center system, but there is no disclosure of synchronous copying from the first site to the second site and asynchronous copying to the third site, and no teaching of using journal data in a journal of the first site. Therefore, it does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

9. U.S. Patent Publication No. 2003/0177321 A1

The published application to Watanabe discloses a technique for enhancing the performance of storage systems to allow recovery after all types of suspensions in remote copy operations. Data is synchronized after an interruption in transfer between a first storage volume of a primary storage system 12 and a first storage volume of a secondary storage system 14 which also includes a second storage volume. After the interruption is detected, at the primary storage system, a record is provided of the data written onto the first storage

volume of the primary storage system, and at the secondary storage volume a record is provided of the data written onto the first storage volume of the secondary storage system. Then, at least a partial copy of the record of the data written onto the first storage volume of the primary storage system is written onto the second storage volume. Using the copy, the first storage volume of the secondary storage system is synchronized with the second storage volume of the secondary storage system.

The reference is directed to synchronizing two or more remote or remote data storage facilities so that they hold and maintain the same data images in the event of a failure of the local storage, but it fails to teach the synchronous copying and asynchronous copying and the use of journal data in a journal of the first storage system. Fig. 6 shows the copying of primary bitmap so that the failure of volume 62 and its resulting loss of bitmap does not preclude resynchronization using the bitmaps. See paragraph [0050]. The bitmap generated during suspension of data transfer is not a journal that includes the control data and journal data as recited in the claims.

Therefore, the reference does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

10. U.S. Patent Application No. 10/832,839, filed April 26, 2004

The application to Watanabe discloses a processing system including a local storage system and a number of remote storage systems that store data mirroring the data maintained by the local storage system. Data that is written, changed, deleted or otherwise modified by the local storage system is periodically sent to the remote storage systems to update the mirroring data maintained by each.

The reference is directed to synchronizing two or more remote or remote data storage facilities so that they hold and maintain the same data images in the event of a failure of the local storage, but it fails to teach the synchronous copying and asynchronous copying and the use of journal data in a journal of the first storage system. Therefore, it does not teach a remote copy system involving three storage systems, wherein the write request from the first host to the first storage system completes after the write data are copied to the second storage system, and wherein the journal data of the journal are received to the third data volume independently from completion of the write request according to information provided in the control data, as recited in independent claim 1. Nor does it teach that the second storage system is configured to synchronously receive and store data corresponding to the write data to the first data volume in the second data volume; and that the third storage system is configured to asynchronously receive and store data corresponding to the journal data of the journal in the third data volume according to information provided in the control data, as recited in independent claim 26.

(f) In view of this petition, the Examiner is respectfully requested to issue a first Office Action at an early date.

Respectfully submitted,



Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
Attachments
RL:rl
60497675 v1